

**Amendments to the Claims:**

This listing of claims will replace all prior versions and listings of claims in the application:

**Listing of Claims:**

1. (Currently amended) A method of storing data in a grid-based data storage system comprising:

defining a plurality of arrays an array of equal capacity data storage blocks within said data storage system, each array having a plurality of rows and units, each defined in terms of a plurality of columns such that each column of said plurality of columns corresponds to a separate storage domain of a plurality of storage domains in said data storage system and each row of said plurality of rows corresponds to against a plurality of data storage blocks with one data storage block of said plurality of storage blocks from each storage domain of said plurality of storage domains at the same physical address, said plurality of rows having a plurality of contiguous data storage blocks in each said separate storage domain; and

defining a plurality of allocating each of the data storage units with at least one data storage unit in each array of said plurality of arrays, each data storage unit of said plurality of data storage units having at least two data storage blocks in different columns of said array of data storage blocks for entirely storing either user data or fault tolerance data;

creating metadata that includes an array pointer that identifies each array of said plurality of arrays; and

~~accessing one data storage unit of said plurality of data storage units using said metadata.~~

2. (Currently amended) The method of claim 1 further comprising: aligning a ~~first row of said array of data storage blocks~~ to a base address ~~in said data storage system that is a multiple of the number of rows in said array of data storage blocks multiplied by said capacity of said data storage blocks~~ of each data storage unit allocated for storing user data with an address defined by a multiple of a size of the storage blocks.

3. (Currently amended) The method of claim 1 ~~where said step of creating metadata further comprises comprising:~~ defining a ~~second~~ pointer that ~~specifies one~~ identifies a selected data storage unit of a plurality of data storage units in one array or said plurality of arrays.

4. (Currently amended) The method of claim 1 wherein the fault tolerance data is characterized as either mirror data or parity data 3 ~~where said step of defining a plurality of data storage units further comprises:~~

~~defining a plurality of data storage units in each array that each comprise a plurality of data blocks containing user data and that further comprise a plurality of data blocks that contain mirror data if the data storage unit format is RAID 1 and that comprise at least one data block that contains first parity data if the data storage unit format is RAID 5 and that contain second parity data if the data storage unit format is RAID 6..~~

5. (Canceled)

6. (Canceled)

7. (Canceled)

8. (Currently amended) The method of claim [[4]] 1 further comprising defining a sparing table that designates at least one data block in each row of said array as a spare data block.

9. (Currently amended) The method of claim 8 further comprising defining said sparing table to contain a plurality of sparing versions that each specify a different storage domain as containing said at least one spare data block in each row ~~and such that the sparing version used for an array depends upon the base address of said array.~~

10. (Canceled)

11. (Canceled)

12. (Canceled)

13. (Currently amended) The method of claim 12 further comprising: creating a metadata entry that specifies said at least one identifies each data storage unit and the array of said a plurality of arrays in which said at least one data storage unit exists.

14. (Currently amended) The method of claim 13 wherein said metadata entry further comprises: an entry indicating the a storage format of said each data storage unit.

15. (Canceled)

16. (Currently amended) A ~~data storage grid for storing data in a grid-based~~ data storage system comprising: an array of equal capacity data storage ~~blocks~~ units, each defined ~~within said data storage system having in terms of~~ a plurality of rows and a plurality of columns such that each column of said plurality of columns corresponds to a separate storage domain of a plurality of storage domains in said data storage system and each row of said plurality of rows corresponds to domains against a plurality of data storage blocks, wherein each data storage unit is individually allocated for entirely storing either user data or fault tolerance data such that one storage block from each column of said plurality of columns has the same physical address, said plurality of rows having a plurality of contiguous data storage blocks in each domain;

~~at least one data storage unit defined in said array that corresponds to a predetermined storage format and that specifies the locations in the array of data blocks containing user data and specifies the locations in the array of data blocks containing mirror data if said predetermined storage format is RAID 1 and~~

~~specifies the location in the array of at least one data block containing row parity if said predetermined storage format is RAID 5 and specifies the location in the array of at least one data block containing vertical parity if said predetermined storage format is RAID 6, said locations in the array of data blocks containing user data organized such that at least one row in said array does not contain mirror data and does not contain row parity and does not contain vertical parity; and a pointer that identifies said at least one data storage unit.~~

17. (Currently amended) The data storage ~~grid system~~ of claim 18 16 further comprising: a plurality of ~~data spare~~ storage blocks ~~designated as spare that are allocatable for storing~~ data storage blocks in said array in the respective data storage units.

18. (Currently amended) A method of ~~converting a data storage grid in a storage system comprising a plurality of data storage grids from a first storage format to a second storage format, said data storage grid having a plurality of data storage blocks configured as a two dimensional array and a predetermined number of data storage units corresponding to said first data storage format forming an ordered set of data storage units defined in said array such that each data storage unit of said plurality of data storage units has a predefined number of user data storage blocks, the method comprising:~~  
~~identifying said data storage grid and the data storage format of said grid;~~  
~~accessing metadata for said first data storage format and for said second data storage format that defines the locations and content of data blocks in a grid and that specifies data blocks employed to calculate row parity if the data storage~~

~~format includes row parity and that specifies the data blocks employed to calculate vertical parity if the data storage format includes vertical parity;~~

~~determining the number of data storage units in said first data storage format;~~

~~determining the number of data storage units in said second data storage format;~~

~~allocating a new grid of the second data storage format if the number of data storage units for said first data storage format exceed the number of data storage formats for said second data storage format;~~

~~copying data storage units in excess of the number of data storage units supported by said second format to said new grid if the number of data storage units for said first data storage format exceeds the number of data storage formats for said second data storage format;~~

~~copying said data storage units in excess of the number of data storage units supported by said second format to said new grid to create mirror data if the second data storage format is RAID 1 and if the number of data storage units for said first data storage format exceeds the number of data storage formats for said second data storage format;~~

~~calculating row parity in said new grid if said second data storage format includes row parity and if the number of data storage units for said first data storage format exceeds the number of data storage formats for said second data storage format;~~

~~calculating vertical parity in said new grid if said second data storage format includes vertical parity and if the number of data storage units for said first data storage format exceeds the number of data storage formats for said second data storage format;~~

~~designating data storage units in said data storage grid copied to said new grid as free if the number of data storage units for said first data storage format exceeds the number of data storage formats for said second data storage format;~~

~~copying data storage units in said data storage grid to create mirror data if said second data storage format is RAID 1;~~

~~calculating and writing row parity in said data storage grid as described by said metadata if said second data storage format is RAID 5;~~

~~calculating and writing row parity in said data storage grid as described by said metadata if said second data storage format is RAID 6; and~~

~~calculating and writing vertical parity in said data storage grid as described by said metadata if said second data storage format is RAID 6~~

storing data in a grid-based data storage system comprising:

defining an array of equal capacity data storage units, each defined in terms of a plurality separate storage domains against a plurality of data storage blocks;

allocating each of the data storage units for entirely storing either user data or fault tolerance data in accordance with a selected first storage format; and

reallocating the data storage units for entirely storing either user data or fault tolerance data in accordance with a selected second storage format if the number of data storage units in the first storage format is less than or equal to the number of data storage units in the second storage format.

19. (Currently amended) The method of claim 18 further comprising:

~~determining the logical device to which said data storage grid is allocated;~~

~~allocating said new data grid to said logical device; and~~

~~updating metadata for said logical device to designate said data storage grid as said~~

~~second data storage format~~

defining a second array of equal capacity data storage units, each defined in terms of

a plurality of separate storage domains against a plurality of data storage blocks;

and

allocating additional data storage units in the second array if the number of data

storage units in the first storage format is greater than the number of data storage

units in the second storage format.

20. (Currently amended) ~~The A method of claim 18 sparing a data storage grid in a~~

~~storage system comprising a plurality of data storage grids in response to a failed storage~~

~~domain, each data storage grid of said plurality of data storage grids having a plurality of~~

~~data storage blocks configured as a two dimensional array having a plurality of columns each~~

~~corresponding to a separate storage domain of a plurality of storage domains in said storage~~

~~system and a plurality of rows each corresponding to a data storage block in each said~~

~~separate domain at the same physical address and having at least one spare data block in each~~

~~row of said plurality of rows, said method comprising:~~

~~identifying said data storage grid in said storage system;~~

~~determining the data storage format of said data storage grid array;~~

~~identifying said an inaccessible failed storage domain;~~

~~accessing a sparing table that specifies a domain for said at least one spare data block~~

~~in each row of said grid array;~~

~~identifying a non-accessible data block in each row of said grid that corresponds to  
said failed storage domain;~~

accessing metadata corresponding to said data storage format that specifies the  
location and content of each data block in said ~~data storage grid array~~; and  
regenerating or copying data corresponding to each ~~said~~ non-accessible data storage  
~~block in each row of said grid and storing said data in said at least one to the  
respective~~ spare data block in each row as identified by said sparing table.

21. (Currently amended) A data structure ~~that manages access to a storage system~~  
~~having a plurality of data grids each having an array of equal capacity data storage units,~~  
~~each defined in terms of blocks within said data storage system, said array having a plurality~~  
~~of rows and a plurality of columns such that each column of said plurality of columns~~  
~~corresponds to a separate storage domain of a plurality of storage domains in said data~~  
~~storage system and each row of said plurality of rows corresponds to against a plurality of~~  
~~data storage blocks, such that one storage block from each column of said plurality of~~  
~~columns has the same physical address, and said plurality of rows having a plurality of~~  
~~contiguous data storage blocks in each domain, said data grids further having a plurality of~~  
~~data storage units each having a predefined group of data storage blocks, said the data~~  
~~structure comprising: a data grid pointer that when multiplied by the a storage capacity of~~  
~~said data storage blocks and by the number of rows in said data storage grid of data storage~~  
~~units allocated for storing user data, and to which an offset, if any, is added, produces the~~  
~~physical base address in each storage domain at which said data storage grid begins;~~

~~a data grid map that specifies the row and column of each data block of said predefined group of data storage blocks for each data storage unit of said plurality of data storage units; and~~

~~a data storage unit pointer selects said row and column of each data block of said predefined group of data blocks for one data storage unit of said plurality of data storage units.~~

22. (Canceled)

23. (Canceled)

24. (Canceled)

25. (Currently amended) A method of accessing data in a data storage system comprising:

defining an array ~~a plurality of data storage grids in said data storage system, each having an array of equal capacity data storage blocks units, each defined in terms of~~ within said data storage system, said array having a plurality of rows and a plurality of columns such that each column of said plurality of columns corresponds to a separate storage domain of a plurality of storage domains in said data storage system and each row of said plurality of rows corresponds to against a plurality of data storage blocks, such that one data storage block from each column of said plurality of columns has the same physical address, and said

~~plurality of rows having a plurality of contiguous data storage blocks in each domain, said data grids further having a plurality of data storage units each having a predefined group of data storage blocks, said array aligned to a base address in said data storage system that is multiple of the number of rows in said array multiplied by said capacity of said data storage blocks;~~

~~creating metadata for said plurality of data grids comprising a data grid pointer, a data storage unit pointer, and a data grid map specifying the row and column in said data grid of data blocks comprising said predefined group of data storage blocks for each of said plurality of data storage units;~~

~~allocating a plurality of said data storage units to a logical device by specifying a data grid pointer value and a data storage unit pointer value for each data storage unit of said plurality of data storage units allocated;~~

~~receiving a request for access to a logical block of said logical device;~~

determining a domain and a physical address for at least one data storage block by applying ~~said data grid~~ an array pointer value to ~~said metadata~~ to determine a base address of ~~said data grid~~ array, applying ~~said a~~ data storage unit pointer value to ~~said data grid map~~ to determine at least one domain to be accessed, and to determine a row such that the row number minus one is multiplied by ~~said a~~ capacity of said storage blocks and added to ~~said a~~ base address to produce a physical address; and

issuing an access command to ~~said~~ at least one domain that includes ~~said~~ physical address.

26. (Canceled)

27. (New) The data storage system of claim 16 wherein the data storage blocks are of equal capacity.

28. (New) The data storage system of claim 16 wherein within a particular row, a data storage block in a first domain has the same physical address as a data storage block in a second domain.

29. (New) The data storage system of claim 16 wherein the array is adapted for storing data in the data storage blocks in a selected storage format by predetermined arrangements of user data, mirror data, and parity data associated with each of a plurality of selectable storage formats.

30. (New) The data storage system of claim 29 wherein the predetermined arrangement defines at least one row in the array that contains no mirror or parity data in the data storage blocks.

31. (New) The data storage system of claim 30 wherein the array is adapted for storing data in a RAID 5 or RAID 6 storage format.

32. (New) The data storage system of claim 16 further comprising metadata defining a pointer for identifying a selected data storage unit.

33. (New) The data storage system of claim 16 wherein a base address of each data storage unit allocated for storing user data is aligned with an address defined by a multiple of a size of the storage blocks.

34. (New) The data storage system of claim 16 wherein the fault tolerance data comprises either mirror data or parity data.